Robust and Parallel Bayesian Model Selection

Michael Zhang, Henry Lam, Lizhen Lin

October 21, 2016

arXiv:1610.06194v1 [stat.ML] 19 Oct 2016

Abstract

Effective and accurate model selection that takes into account model uncertainty is an important but challenging problem in modern data analysis. One of the major challenges is the computational burden required to infer huge data sets which, in general, cannot be stored or processed on one machine. Moreover, in many real data modeling scenarios we may encounter the presence of outliers and contaminations that will damage the quality of our model and variable selection.

We can overcome both of these problems through a simple "divide and conquer" strategy in which we divide the observations of the full data set equally into subsets and perform inference and model selections independently on each subset. After local subset inference, we can aggregate the optimal subset model or aggregate the local model/variable selection criteria to obtain a final model. We show that by aggregating with the geometric median, we obtain results that are robust to outliers and contamination of an unknown nature.

1 INTRODUCTION

In many data modeling scenarios, a large number of plausible models are available to fit to the data, each of which may result in drastically differing predictions and conclusions. Being able to select the right model or to properly incorporate model uncertainty for inference is an crucial task. Here we consider model selection for a normal linear model:

$$Y = X\beta + \epsilon, \;\; \epsilon \sim N(0, \sigma^2 I), \tag{1}$$

where Y is an N dimensional response vector, X is an $N \times D$ dimensional design matrix and β is a D dimensional vector of regression parameters. In a Bayesian setting, we have a natural probabilistic evaluation of models through the posterior model probability. Depending on the objectives of the data analysis, we may be interested in selecting the "best" model or obtaining predictions with minimum error.

Existing procedures to accomplish the aforementioned goals, however, will perform poorly under the presence of outliers and contaminations. In addition, Markov chain Monte Carlo (MCMC) algorithms for these methods do not scale to "big data" situations. To overcome these problems, we introduce a "divide and conquer" method and combine it with existing Bayesian model selection techniques for normal linear models that is robust to outliers which, by consequence, will allow us to perform Bayesian model selection *in parallel*, thus also providing major reductions in computational burden with big data.

Our robust parallelization strategy is based on the ideas for robust posterior inference in [Minsker et al., 2014]. In particular, the notion of the geometric median and the median posterior has shown to be effective in robust and parallel statistical inference [Wang et al., 2014, Minsker, 2015]. Previous work in this area has focused on the performance in terms of the errors relative to the true distribution. Our novel contribution in this paper is to demonstrate the effectiveness of the divide-and-conquer strategy in selecting the correct *class* of models under the domain of Bayesian model selection.

We theoretically quantify the robustness of our strategy and, while doing so, we also show that the model aggregated across different subsets (the "divide") has improved concentration to the true model class compared to the one using the full data set. We conduct numerical studies on simulation data to show that our procedure yields robust predictions results and effectively *selects the correct model* in the presence of outliers and contaminations.

2 BAYESIAN MODEL SELECTION

In Bayesian model selection, we define the prior model probability $Pr(M_k)$ for each of the model M_k $(k = 1, \ldots, K)$ under consideration. For model M_k, we additionally have parameters (β_k, σ_k^2) with prior $Pr(\beta_k, \sigma_k^2|M_k)$, which leads to a likelihood $Pr(Y|\beta_k, \sigma_k^2, M_k)$. Thus, the posterior model probability for model M_k, $Pr(M_k|-)$, is proportional to

$$Pr(M_k) \int Pr(Y|\beta_k, \sigma_k^2, M_k) Pr(\beta_k, \sigma_k^2|M_k)\, \mathrm{d}\beta_k \mathrm{d}\sigma_k^2$$

However, as noted in [Barbieri and Berger, 2004], choosing the model with the highest posterior model probability is not always the best option nor should one neglect the risk of model uncertainty. Instead of resorting to a single model for predicted values $\tilde{Y}$ (or some quantity of interest in general), [Hoeting et al., 1999] proposes to average over the model uncertainty with Bayesian model averaging (BMA) to obtain a posterior mean and variance of $\tilde{Y}$ at a covariate level $\tilde{X}$:

$$\begin{aligned} E[\tilde{Y}|\tilde{X}, Y] &= \sum_{k=1}^{K} E[\tilde{Y}|\tilde{X}, Y, M_k] Pr(M_k|\tilde{X}, Y), \\ Var(\tilde{Y}|\tilde{X}, Y) &= \sum_{k=1}^{K} Pr(M_k|X, Y) \Big(Var(\tilde{Y}|\tilde{X}, Y, M_k) + \\ &\qquad E[\tilde{Y}|\tilde{X}, Y, M_k]^2 \Big) - E[\tilde{Y}|X, Y]^2. \end{aligned}$$

We will focus on BMA in our theoretical developments in this paper. Our numerical experiments, however, will show that our divide-and-conquer strategy is also effective in applying on other model selection methods.

The first alternative to BMA is the median probability model, which can be shown to be optimal if we must choose one model for prediction [Barbieri and Berger, 2004]. First, define the posterior inclusion probability of each predictor x_d $(d = 1, \ldots, D)$ as the sum of posterior model probabilities of the models that include predictor x_d, namely $p_d =$

$\sum_{k:x_d \in M_k} Pr(M_k|X,Y)$. The median probability model is the model that includes the predictors x_d if $p_d \geq 1/2$.

Second, using the maximum value of the likelihood for each model $Pr(Y|\hat{\beta}_k, \hat{\sigma}_k^2, M_k)$, where $(\hat{\beta}_k, \hat{\sigma}_k^2)$ is the maximum likelihood estimate of (β_k, σ_k^2), we can perform penalized model selection through the Akaike information criterion (AIC) [Akaike, 1974] or the Bayesian information criterion (BIC) [Schwarz, 1978] by selecting the model with the lowest information criterion.

$$\begin{aligned} \text{AIC} &= -2\log Pr(Y|\hat{\beta}_k, \hat{\sigma}_k^2, M_k) + 2(D+1) \\ \text{BIC} &= -2\log Pr(Y|\hat{\beta}_k, \hat{\sigma}_k^2, M_k) + (D+1)\log N. \end{aligned}$$

The final model selection technique we will consider is stochastic variable selection through the spike and slab model [George and McCulloch, 1993], which allows for variable shrinkage under high-dimensional models. For the purposes of this paper, we will use the rescaled spike and slab model [Ishwaran and Rao, 2005]. To perform posterior inference in the rescaled spike and slab model, first define $Y' = \sqrt{\frac{N}{\hat{\sigma}^2}}Y$ where $\hat{\sigma}^2$ is the unbiased estimate of σ^2 under the full model and let $\nu_0 > 0$ be some small number. The model is defined to be following mixture model:

$$\begin{aligned} Y' &\sim N(X\beta, N\sigma^2 I),\ \beta_{SS_d} \sim N(0, J_d\tau_d^2), \\ \sigma_{ss}^{-2} &\sim \text{Gamma}(a,b),\ J_d \sim (1-w)\delta_{J_d}(\nu_0) + w\delta_{J_d}(1), \\ \tau_d^{-2} &\sim \text{Gamma}(a_\tau, b_\tau),\ w \sim \text{Uniform}(0,1). \end{aligned}$$

3 DIVIDE-AND-CONQUER AND ROBUST BAYESIAN MODEL SELECTION

In our robust model selection strategy, we divide N observations into R subsets of roughly equal sample size. Then inference, model selection and prediction is performed for the linear model $Y_{(j)} = X_{(j)}\beta + \epsilon_{(j)}$ independently across $j = 1, \ldots, R$ subsets using the existing Bayesian model selection procedures, which are then combined to form a final model or a combined prediction value.

Given linear model (1), we first define the following priors on a normal likelihood with response variable Y and D-dimensional predictor X. The N observations are divided into R subsets with N_j observations within subsets. One has,

$$\begin{aligned} Pr\left(\sigma^{-2}\right) &\sim \text{Gamma}(a,b), \\ Pr(\beta|\sigma^2) &\sim N(\beta_0, \sigma^2\Sigma_0). \end{aligned}$$

To compensate for the data division, we raise the likelihood of the divided data $Pr(Y_{(j)}|X_{(j)}, \beta, \sigma^2)$ to the R-th power and adjust the normalizing constant accordingly so that the likelihood for Y_j is:

$$\left(\frac{R}{2\pi\sigma^2}\right)^{N/2} \exp\left\{\frac{-R}{2\sigma^2}\left(Y - X\beta\right)^T\left(Y - X\beta\right)\right\}.$$

The intuition and motivation for raising the subset likelihood to R-th power is to adjust the potentially inflated variance of the subset posterior distribution. Exploiting conjugacy, we obtain the full conditionals:

$$\begin{aligned}
Pr(\beta|-) &\sim N\left(\mu_\beta, \sigma^2\Sigma_\beta\right), \\
\mu_\beta &= \Sigma_\beta\left(\beta_0\Sigma_0^{-1} + RX^TY\right), \\
\Sigma_\beta &= \left(\Sigma_0^{-1} + RX^TX\right)^{-1}, \\
Pr\left(\sigma^{-2}|-\right) &\sim \text{Gamma}\left(a', b'\right), \\
a' &= a + \frac{N+D}{2}, \\
b' &= b + \frac{R}{2}\epsilon^T\epsilon + \frac{1}{2}\left(\beta-\beta_0\right)^T\Sigma_0^{-1}\left(\beta-\beta_0\right), \\
\epsilon &= \left(Y - X\beta\right).
\end{aligned}$$

Let $\Sigma_X = I + RX\Sigma_0X^T$, then integrating out the parameters gives us the following marginal distribution $Pr(Y_{(j)}|X_{(j)})$:

$$\frac{\left(\frac{R}{2\pi}\right)^{\frac{N}{2}} b^a\Gamma(a+\frac{N}{2})\left|\Sigma_X\right|^{-\frac{1}{2}} / \Gamma(a)}{\left(b + \frac{R}{2}\left(Y - X\beta_0\right)^T\Sigma_X^{-1}\left(Y - X\beta_0\right)\right)^{a+\frac{N}{2}}}.$$

For distributed AIC and BIC model evaluation, we raise the likelihood term of the AIC and BIC formula to the power of R:

$$\begin{aligned}
\text{AIC}_R &= -2R\log Pr(Y|\hat{\beta}_k, \hat{\sigma}_k^2, M_k) + 2(D+1) \\
\text{BIC}_R &= -2R\log Pr(Y|\hat{\beta}_k, \hat{\sigma}_k^2, M_k) + (D+1)\log N.
\end{aligned}$$

In applying our procedure with the spike and slab prior, we derived the full Gibbs sampler for our procedure. The details are included in the supplementary materials.

Once inference is built on each subset, the key step is to aggregate the subset models (or estimates) together into a final model (or estimate). To aggregate our results, we collect the R number of subset models or estimates and find the geometric median between these R elements. The geometric median for a set of elements $\{x_1, \ldots, x_R\}$ valued on a Hilbert space $\mathbb{H}$, is defined as

$$\begin{aligned}
x_* &= \text{med}_g(x_1, \ldots, x_R) \\
&= \text{argmin}_{y\in\mathbb{H}}\sum_{j=1}^{R}\|y - x_j\|
\end{aligned} \tag{2}$$

where $\|\cdot\|$ is the norm associated with the inner product in $\mathbb{H}$ [Minsker et al., 2014]. The solution can generally be effectively approximated using the Weiszfeld algorithm [Weiszfeld, 1937].

For instance, in the case of aggregating the posterior model probabilities across R subsets of data, the geometric median operates on the space of posterior distributions and the

geometric median posterior model probability, $Pr_*(M_k|X,Y)$, is defined as:

$$\underset{P\in\Pi_K}{\operatorname{argmin}} \sum_{j=1}^{R} \big|\big|P - Pr(M_k|X_{(j)},Y_{(j)})\big|\big| \tag{3}$$

where $Pr(M_k|X_{(j)},Y_{(j)})$ is the posterior model probabilities for subset j, and Π_K denotes the space of distributions on K support points. The metric $\|\cdot\|$ here can be taken as the Euclidean metric, or an integral probability metric (IPM) defined as $||P-Q|| = \sup_{f\in\mathcal{F}} \left|\int f(x)\,\mathrm{d}(P-Q)(x)\right|$ for some class of functions $\mathcal{F}$ [Sriperumbudur et al., 2010, Sriperumbudur et al

For the model selection techniques discussed earlier (AIC, BIC, and the median model selection), we can choose a final model in two ways: One, we can select the best model locally on each subset, use it for prediction, and then aggregate the results (estimate combination). Or two, we can take the median of the model selection criteria and choose that particular model on each subset and then aggregate the results to get a final model (model combination).

However, in Bayesian model averaging and spike and slab modeling we do not choose a final model. We can still perform model or estimate combination by aggregating the posterior model probabilities. We consider both model and estimate combinations in our experiments (in the numerical example later and also in the Supplementary Materials) and show that they yield similar results.

Algorithm 1: Algorithm for robust model selection in the case of BMA.

for $j \in \{1,\ldots,R\}$ **do**
- Raise likelihood to R-th power
- Compute inference for $P(\theta|M_k,X_{(j)},Y_{(j)})$ for $k=1,\ldots,K$
- Draw predictive values from predictive posterior $P(\tilde{Y}|M_k,X_{(j)},Y_{(j)})$ for $k=1,\ldots,K$
- Calculate posterior model probabilities $\{P(M_k|X_{(j)},Y_{(j)})\}_{k=1,\ldots,K}$

Calculate geometric median of posterior model probabilities over the subsets using (3).
Approximate geometric medians of posterior parameter probabilities or predictive values given individual models over the subsets using (2).
Obtain BMA estimate: $E[\tilde{Y}|Y,X] = \sum_{k=1}^{K} E_*[\tilde{Y}|X,Y,M_k]Pr_*(M_k|X,Y)$

4 THEORETICAL ROBUSTNESS AND IMPROVED CONCENTRATION

In this section we provide theoretical justification on the robustness in the divide-and-conquer strategy. While we focus on BMA, our arguments naturally extend to other model selection techniques.

In addition to robustness, we also show, in a well-defined sense, that the aggregated model class is closer to the true one in terms of concentration compared to using the whole data set at once.

The proofs of all theorems in this section are presented in the Supplementary Materials. Note that we do not raise the subset likelihood to R-th power in our current theoretical analysis, but the results can be generalized by imposing slightly stronger entropy conditions on the model.

Let $\mathcal{S}$ be the domain of $\theta = (M_k, \beta, \sigma^2)$, our set of model indices and parameters. Let θ_0 be the truth, and let (X_1, Y_1) be a generic data point. Let $p_0(y|x) := p(y|x, \theta_0)$ be the true conditional density of Y_1 given X_1, and $p_0(x)$ be the true density of the covariates X_1. We denote $p_\theta(y|x) := p(y|x, \theta)$.

Let P_θ be the distribution defined by $p_0(x) \times p_\theta(y|x)$ and P_0 is the true distribution $p_0(x) \times p_0(y|x)$. For convenience, we denote $P_0 f = P_0 f(X_1, Y_1) = E_{p_0}[f(X_1, Y_1)]$ where $E_{p_0}[\cdot]$ is the expectation under $p_0(y|x) \times p_0(x)$. We denote P_0^N as the true probability measure taken on the data (X, Y) of size N and $P_0^N f = E_{P_0^N}[f(X, Y)]$.

Lastly, we denote $\mathcal{D}(\epsilon, \mathcal{P}, d)$ as the ϵ-packing number of a set of probability measures $\mathcal{P}$ under the metric d, which is the maximal number of points in $\mathcal{P}$ such that the distance between any pair is at least ϵ. We implicitly assume here that $\mathcal{P}$ is separable.

The following Theorem 1 follows from a modification of Theorem 2.1 in [Ghosal et al., 2000]:

Theorem 1 *Assume that there is a sequence ε_N such that $\varepsilon_N \to 0$ and $N\varepsilon_N^2 \to \infty$ as $N \to \infty$, a constant C, and a set $\mathcal{S}_N \in \mathcal{S}$ so that*

1. $\log \mathcal{D}(\varepsilon_N/2, \mathcal{P}_{\mathcal{S}_N}, d_H) \leq N\varepsilon_N^2$
2. $Pr(\mathcal{S} \setminus \mathcal{S}_N) \leq e^{-N\varepsilon_N^2(C+4)}$
3. $Pr\left(\theta : -P_0 \log \frac{p_\theta(Y_1|X_1)}{p_0(Y_1|X_1)} \leq \varepsilon_N^2, \right.$
   $\left. P_0 \left(\frac{p_\theta(Y_1|X_1)}{p_0(Y_1|X_1)}\right)^2 \leq \varepsilon_N^2 \right) \geq e^{-N\varepsilon_N^2 C}$

where $\mathcal{P}_{\mathcal{S}_N} = \{p_0(x) \times p_\theta(y|x) : \theta \in \mathcal{S}_N\}$ and d_H is the Hellinger distance. Then we have

$$P_0^N \left(Pr(\theta : d_H(P_\theta, P_0) > T\varepsilon_N^2 | X, Y) > \delta\right) \leq \frac{1}{C^2 N\varepsilon_N^2 \delta} + \frac{2e^{-LN\varepsilon_N^2}}{\delta} + \frac{2e^{-2N\varepsilon_N^2}}{\delta} \tag{4}$$

for any $0 < \delta < 1$ and sufficiently large $T > 0$ such that $LT^2 \geq C + 4$ and $LT^2 - 1 > L$, where L is a universal constant.

As noted by [Ghosal et al., 2000], the important assumptions are Assumptions 1 and 3. Essentially, Assumption 1 constrains the size of the parameter domain $\mathcal{S}$ to be not too big, whereas Assumption 3 ensures sufficient mass of the prior on a neighborhood of the true parameter.

The concentration result (4) states that the posterior distribution of θ is close to the true θ_0 with high probability, where the closeness is measured in terms of the Hellinger distance between the likelihoods. Note that the RHS of (4) consists of three terms. The dominant term is the power-law decay in $N\varepsilon_N^2$.

The other two exponential decay terms result from technical arguments in the existence of tests that sufficiently distinguish between distributions [Birgé, 1983, Le Cam, 2012].

Next we describe the concentration behavior of BMA. We focus on the situations where all the candidate models are non-nested, i.e. only one model contains distributions that are arbitrarily close to the truth. Without loss of generality, we let M_1 be the true model. Our result can be generalized to the nested situation, where the concentration statement is modified for the probability of selecting one of the correct models.

Theorem 2 (BMA of Non-Nested Models) *Suppose the assumptions in Theorem 1 hold. Also assume that, for sufficiently small $\epsilon > 0$, $d(P_\theta, P_0) > \epsilon$ for any $\theta \in \mathcal{S}_{-1} := \{\theta = (M_k, \beta, \sigma^2) : k \neq 1\}$. Let T be chosen as in Theorem 1. We have*

1. *For any $0 < \delta < 1$,*

$$P_0^N \left(Pr(M_1|X,Y) < 1-\delta\right) \leq \frac{1}{C^2 N \varepsilon_N^2 \delta} + \frac{2e^{-LN\varepsilon_N^2}}{\delta} + \frac{2e^{-2N\varepsilon_N^2}}{\delta}. \tag{5}$$

2. *For any $0 < \delta < 1$,*

$$P_0^N \left(d_E(Pr(M_k|X,Y), \mathbf{e}_1) > \delta\right) \leq \frac{\sqrt{2}}{C^2 N \varepsilon_N^2 \delta} + \frac{2\sqrt{2}e^{-LN\varepsilon_N^2}}{\delta} + \frac{2\sqrt{2}e^{-2N\varepsilon_N^2}}{\delta}, \tag{6}$$

   where d_E is the Euclidean distance, and $\mathbf{e}_1$ is the point mass on M_1.

3. *For any $0 < \delta < \sqrt{(\sqrt{2}-1)^2+1}/2$,*

$$P_0^N \left(d_H(Pr(M_k|X,Y), \mathbf{e}_1) > \delta\right) \leq \frac{(\sqrt{2}-1)^2+1}{\sqrt{2}C^2 N \varepsilon_N^2 \delta^2} + \frac{((\sqrt{2}-1)^2+1)e^{-LN\varepsilon_N^2}}{\delta^2} + \frac{((\sqrt{2}-1)^2+1)e^{-2N\varepsilon_N^2}}{\delta^2}. \tag{7}$$

Note that the assumption $d(P_\theta, P_0) > \epsilon$ for any $\theta \in \mathcal{S}_{-1}$ and sufficiently small ϵ is a manifestation of the non-nested model situation, asserting that only one model is "correct". Result 1 is a concentration on the posterior probability of picking the correct model to be close to 1.

Result 2 translates this in terms of the Euclidean distance between the model posterior probability and the point mass on the correct model. Result 3 is an alternative using the Hellinger distance. In addition to the proof, we also provide some explanation of the inferior concentration bound for Hellinger distance (7) compared to Euclidean distance (6) in the Supplementary Materials.

The following result states how a divide-and-conquer strategy can improve the concentration rate of the posterior model probabilities towards the correct model:

Theorem 3 (Concentration Improvement) *Suppose the assumptions in Theorem 2 hold. Let* $s = N/R$, *and* $q = \frac{\sqrt{2}}{C^2 N \varepsilon_N^2 \delta} + \frac{2\sqrt{2}e^{-LN\varepsilon_N^2}}{\delta} + \frac{2\sqrt{2}e^{-2N\varepsilon_N^2}}{\delta}$ *For sufficiently large* N, *letting* α, ν *be constants such that* $0 < q < \alpha < 1/2$ *and* $0 \leq \nu < (\alpha - q)/(1 - q)$, *we have:*

1. $Pr_*(M_k|X, Y)$, *the geometric median under* d_E *of* $\{Pr(M_k|(X_{(j)}, Y_{(j)}))\}_{j=1,\ldots,R}$, *satisfies*

$$P_0^N \left(d_E(Pr_*(M_k|X, Y), \mathbf{e}_1) > C_\alpha \delta \right) \leq \left(e^{(1-\nu)\psi(\frac{\alpha-\nu}{1-\nu}, q)} \right)^{-R} \tag{8}$$

   where $C_\alpha = (1 - \alpha)\sqrt{1/(1 - 2\alpha)}$, *and* $\psi(\alpha, q) = (1 - \alpha) \log \frac{1-\alpha}{1-q} + \alpha \log \frac{\alpha}{q}$.

2. *Let* K *be the number of model classes, then:*

$$P_0^N \left(Pr_*(M_1|X, Y) < 1 - C_\alpha \delta \sqrt{\frac{K-1}{K}} \right) \leq \left(e^{(1-\nu)\psi(\frac{\alpha-\nu}{1-\nu}, q)} \right)^{-R} \tag{9}$$

3. *Suppose in addition that, for any* $P_{\theta^1}, P_{\theta^2}$ *such that* $\theta^i = (M_1, \beta^i, (\sigma^2)^i)$ *for* $i = 1, 2$, *we have*

$$d_H(P_{\theta^1}, P_{\theta^2}) \geq \tilde{C} d_{\mathcal{F}}(P_{\theta^1}, P_{\theta^2})^\gamma \tag{10}$$

   for an IPM $d_{\mathcal{F}}$ *associated with the class of functions* $\mathcal{F}$, *and constant* $\tilde{C} > 0$ *and* $\gamma > 0$. *Then*

$$P_0^N \left(Pr_*(M_1|X, Y) > 1 - C_\alpha \delta \sqrt{\frac{k-1}{k}}, \right.$$
$$Pr_* \left(\theta : d_k(P_\theta, P_0) \leq \left(\frac{M\varepsilon_N^2}{\tilde{C}} \right)^{-\gamma} \middle| M_1, X, Y \right)$$
$$\left. \geq 1 - \delta \right) \geq 1 - 2 \left(e^{(1-\nu)\psi\left(\frac{\alpha-\nu}{1-\nu}, q\right)} \right)^{-R}$$

   where $Pr_*(\cdot|M_1, X, Y)$ *is the geometric median over the* R *subsets under* $d_{\mathcal{F}}$ *of the parameter posteriors* $Pr(\cdot|M_1, X_{(j)}, Y_{(j)})$ *conditional on the correct model* M_1.

The significance of Theorem 3 is the improvement of the concentration from power-law decay in Theorem 2 to exponential decay, as the number of subsets grows. Results 1 and 2 describes the exponential concentration for the model posteriors to the correct model, while Result 3 states the joint concentration in both the model posterior and the parameter posterior given the correct model, when one adopts a second layer of divide-and-conquer on the parameter posterior conditional on each individual candidate model. A similar condition as (10) has been used in the context of reproducing kernel Hilbert space [Minsker et al., 2014].

Finally, the following theorem highlights that the concentration improvement still holds even if the data are contaminated to a certain extent:

Theorem 4 (Robustness to Outliers) *Using the notation in Theorem 3, but assume instead that, for j where $1 \leq j \leq \lfloor (1-\nu)R \rfloor + 1$,*

$$P_0^s \left(d_E(Pr(M_k|X_{(j)}, Y_{(j)}), \mathbf{e}_1) > \delta \right) \leq \frac{\sqrt{2}}{C^2 s \varepsilon_s^2 \delta} + \frac{2\sqrt{2}e^{-Ls\varepsilon_s^2}}{\delta} + \frac{2\sqrt{2}e^{-2s\varepsilon_s^2}}{\delta},$$

the conclusion of Theorem 3 still holds.

Theorem 4 stipulates that when a small number of subsets are contaminated by arbitrary nature, the geometric median approach still retains the same exponential concentration. Therefore, our procedure has provable robustness.

5 SIMULATIONS AND DATA ANALYSIS

For the BMA, AIC, BIC and median probability model tests, we generate data from a model $Y = X\beta + \epsilon$, where X is a 5000×10 matrix and β is a 10 dimensional vector with 3 true predictors. We assess the aforementioned model selection techniques with four tests, over 10 trials for the contamination and magnitude tests and over 20 trials for the coverage test on 1 and 10 subsets with 1000 iterations on each MCMC chain and a burn-in period of the initial 500 iterations.

The first test is the contamination test which examines the root mean square error (RMSE) of held-out test data $\tilde{Y}$ of size 50 against the number of outliers present (as many as 5 in our experiments) in the training data, Y.

We generate outliers by taking the maximum of the absolute value of the data and add a given magnitude value. Each outlier has a relative magnitude of 10000. For the contamination test, we expect to see superior performance with regards to RMSE of the 10 subset median posterior as long as the number of outliers per subset does not exceed 1.

Figure 1 demonstrates the robustness of our technique to the number of outliers when we divide the data into subsets. We can see that the empirical 95% distribution of the RMSE for 10 subsets (green dashed line) falls dramatically below that of the RMSE distribution of 1 subset for each model selection technique when outliers are present.

The second test assesses the RMSE of the held-out test data of size 50 against the *increasing relative magnitude of one outlier* present in the training data. We expect to see nearly constant RMSE on the 10 subset run as the relative magnitude of a single outlier increases, thus the procedure is robust.

We can see in Figure 2 that the RMSE of distributed variants of the model selection techniques are lower than the single processor variants as the number of outliers increases. In the magnitude test, we can categorically observe that 10 subset RMSE is invariant to the relative magnitude of one outlier present in the data whereas the RMSE grows rapidly on one subset.

The next test assesses the empirical 95% posterior coverage of the true held-out predictive value of size 1, $\tilde{Y}$, against the *increasing relative magnitude* of one outlier in the training data. To calculate coverage we generate 50 independent MCMC chains at each level of

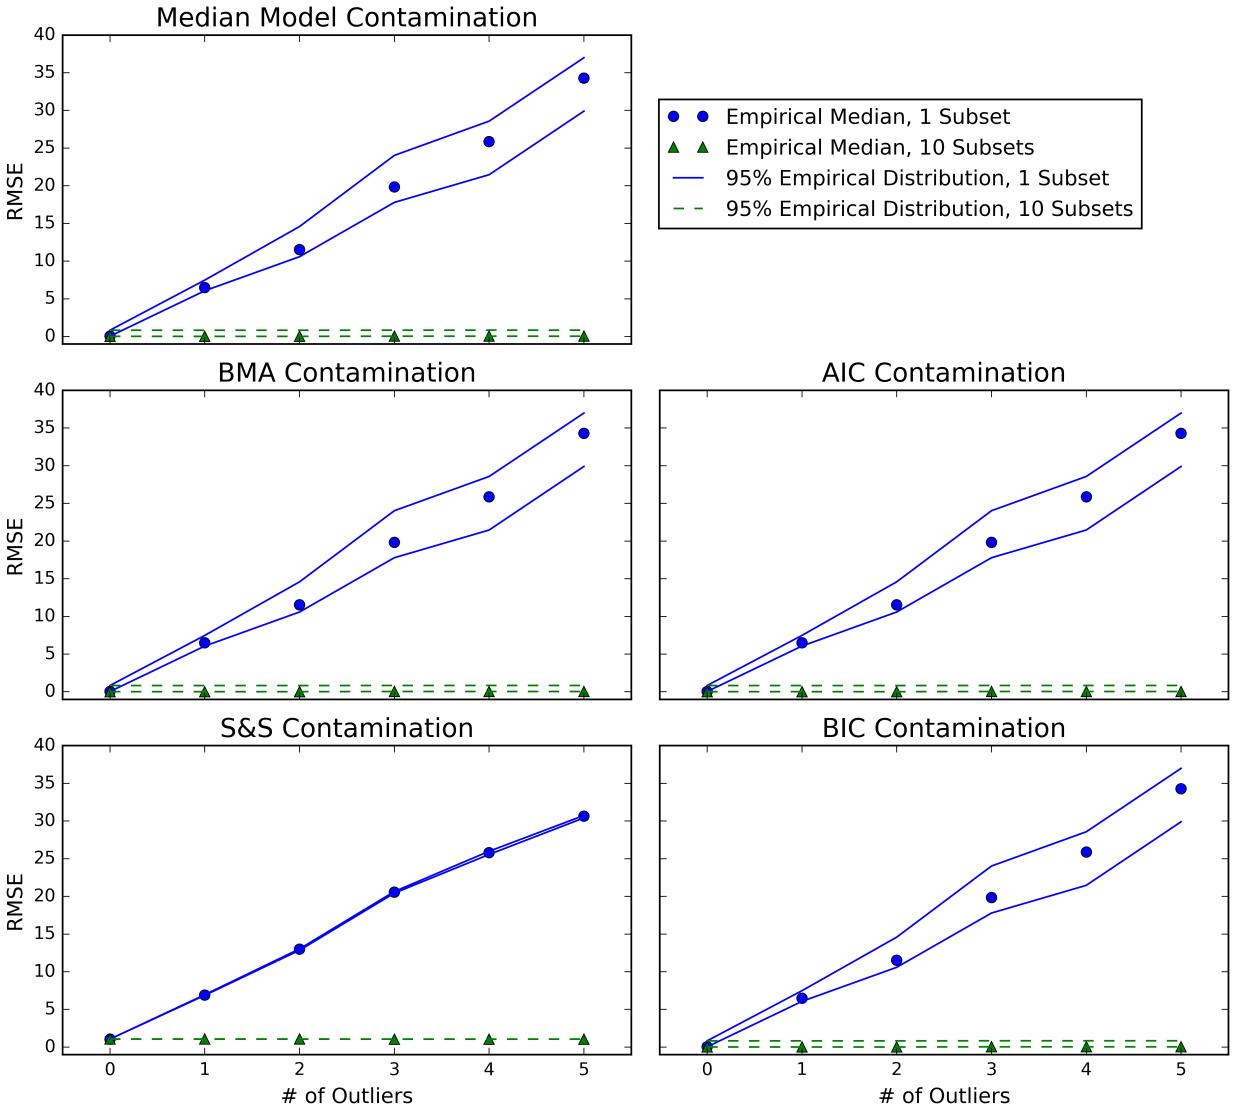

Figure 1: Contamination test.

outlier magnitude and calculate the proportions of chains which include the true predictive value within the 2.5% and 97.5% percentiles of the posterior predictive draws.

For the coverage test we see that the empirical coverage of a single predictive value for the distributed subsets is, on average, 95% regardless of the magnitude of the outlier as opposed to the empirical coverage for the single subset. In the 1 subset case, we can see that the empirical coverage degrades as the magnitude of the outlier grows. (see Figure 3).

Our last evaluation is the coverage of the regression coefficients and the ability for our model selection techniques to *choose the correct model* under the distributed setting with a single outlier of magnitude 10000. We compare the posterior credible interval of the regression coefficients for 1 and 10 subsets. Note that we do not include nested models in our evaluations or models larger than the true model (i.e models with more than 3 covariates included). As seen in Figure 4, the parallel technique is able to select the correct model 1

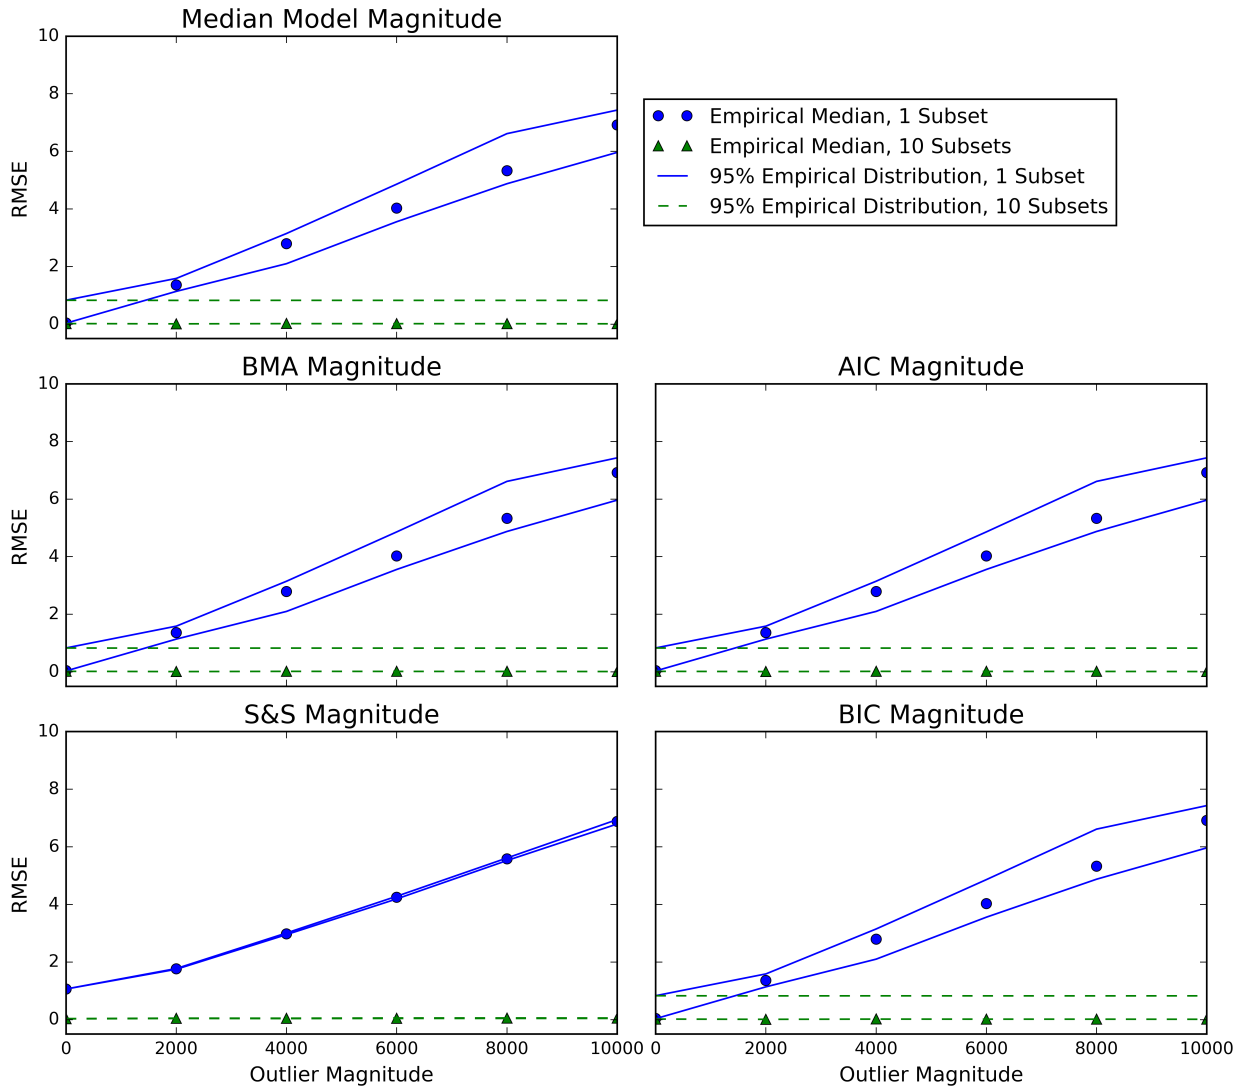

Figure 2: Magnitude of outlier test.

subset test, the outlier leads to the incorrect model being selected. Additionally, Figure 5 demonstrates that model and estimate combination yield similar results with the regression coefficient coverage test. The comparisons between model and estimate combinations for all other tests are contained in the Supplementary Materials.

6 CONCLUSION

Bayesian methods now need to be parallel, scalable and robust for inference in the age of big data. A substantial body of work exists for fast and scalable Bayesian inference methods. But so far, available research methods are lacking on scalable model selection, particularly in the Bayesian framework. We have presented in this paper a technique that allows for the parallelization of an vital task in Bayesian data analysis and machine learning.

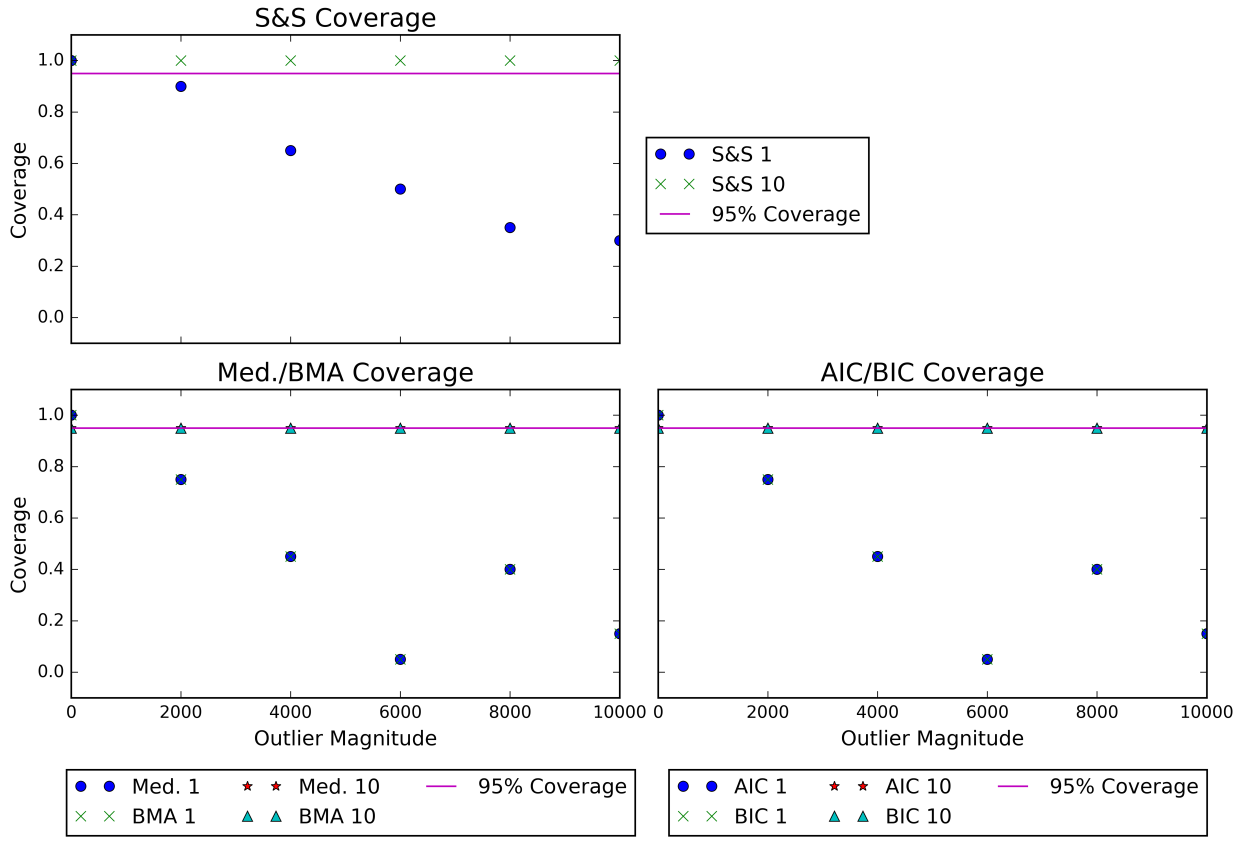

Figure 3: Testing empirical coverage of predictive value.

Our method is simple to execute in practice and requires no communication between machines until the aggregation step after local inference. We show that, through parallelization and proper combination of models, inference is robust to the presence of outliers in the data. Though the selection of model selection criteria and regularization techniques in this paper is not exhaustive, we have shown that our divide and conquer strategy is applicable to a variety of methods for Bayesian model selection.

Acknowledgements

The contribution of LL was funded by NSF grant IIS 1546331.

References

[Akaike, 1974] Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723.

[Barbieri and Berger, 2004] Barbieri, M. M. and Berger, J. O. (2004). Optimal predictive model selection. *Annals of Statistics*, pages 870–897.

[Birgé, 1983] Birgé, L. (1983). Approximation dans les espaces métriques et théorie de l'estimation. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 65(2):181–237.

[George and McCulloch, 1993] George, E. I. and McCulloch, R. E. (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88(423):881–889.

[Ghosal et al., 2000] Ghosal, S., Ghosh, J. K., and van der Vaart, A. W. (2000). Convergence rates of posterior distributions. *The Annals of Statistics*, 28(2):500–531.

[Hoeting et al., 1999] Hoeting, J. A., Madigan, D., Raftery, A. E., and Volinsky, C. T. (1999). Bayesian model averaging: a tutorial. *Statistical science*, pages 382–401.

[Ishwaran and Rao, 2005] Ishwaran, H. and Rao, J. S. (2005). Spike and slab variable selection: Frequentist and Bayesian strategies. *Annals of Statistics*, pages 730–773.

[Le Cam, 2012] Le Cam, L. (2012). *Asymptotic Methods in Statistical Decision Theory*. Springer Science & Business Media.

[Minsker, 2015] Minsker, S. (2015). Geometric median and robust estimation in Banach spaces. *Bernoulli*, 21(4):2308–2335.

[Minsker et al., 2014] Minsker, S., Srivastava, S., Lin, L., and Dunson, D. B. (2014). Scalable and robust Bayesian inference via the median posterior. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 1656–1664.

[Schwarz, 1978] Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464.

[Sriperumbudur et al., 2012] Sriperumbudur, B. K., Fukumizu, K., Gretton, A., Schölkopf, B., Lanckriet, G. R. G., et al. (2012). On the empirical estimation of integral probability metrics. *Electronic Journal of Statistics*, 6:1550–1599.

[Sriperumbudur et al., 2010] Sriperumbudur, B. K., Gretton, A., Fukumizu, K., Schölkopf, B., and Lanckriet, G. R. G. (2010). Hilbert space embeddings and metrics on probability measures. *The Journal of Machine Learning Research*, 11:1517–1561.

[Wang et al., 2014] Wang, X., Peng, P., and Dunson, D. B. (2014). Median selection subset aggregation for parallel inference. In *Advances in Neural Information Processing Systems*, pages 2195–2203.

[Weiszfeld, 1937] Weiszfeld, E. (1937). Sur le point pour lequel la somme des distances de n points donnes est minimum. *Tohoku Mathematical Journal*, 43:355–386.

7 SUPPLEMENTARY MATERIAL

Proofs

Proof of Theorem 1

The proof is a modification of that for Theorem 2.1 in [Ghosal et al., 2000]. Take any $\epsilon > 2\varepsilon_N$, we have, by Assumption 1,

$$\log \mathcal{D}\left(\frac{\epsilon}{2}, \mathcal{P}_{\mathcal{S}_N}, d\right) \leq \log \mathcal{D}\left(\varepsilon_N, \mathcal{P}_{\mathcal{S}_N}, d\right) \leq N\varepsilon_N^2$$

Then, by Theorem 7.1 in [Ghosal et al., 2000], there exists tests ϕ_N and a large enough constant T (chosen later) such that

$$P_0^N \phi_N \leq e^{N\varepsilon_N^2} e^{-LNT^2\varepsilon_N^2} \frac{1}{1 - e^{-LNT^2\varepsilon_N^2}} \tag{11}$$

and

$$\sup_{\theta \in \mathcal{S}_N : d(P_\theta, P_0) > T\varepsilon_N} P_\theta^N (1 - \phi_N) \leq e^{-LNT^2\varepsilon_N^2} \tag{12}$$

for a universal constant $L > 0$, any $N > 0$, and P_θ^N denotes the probability measure on (X, Y) under $(X_1, Y_1) \sim p_0(x) \times p_\theta(y|x)$.

By (11), we have

$$P_0^N Pr(\theta : \theta : d(P_\theta, P_0) > L\varepsilon_N^2 | X, Y)\phi_N \leq P_0^N \phi_N \leq 2e^{-LN\varepsilon_N^2} \tag{13}$$

as $N \to \infty$, if we choose $LT^2 - 1 > L$. Now, since

$$P_0 \frac{p_\theta(Y_1|X_1)}{p_0(Y|X)} = \int \frac{p_\theta(y|x)}{p_0(y|x)} p_0(dy|x) p_0(dx) = \int p_\theta(dy|x) p_0(dx) = 1$$

by Fubini's theorem, we have

$$P_0^N \int_{\mathcal{S} \setminus \mathcal{S}_N} \prod_{i=1}^N \frac{p_\theta(Y_i|X_i)}{p_0(Y_i|X_i)} Pr(\mathrm{d}\theta) \leq Pr(\mathcal{S} \setminus \mathcal{S}_N)$$

Hence, by Fubini's theorem again,

$$\begin{aligned}
& P_0^N \int_{\theta \in \mathcal{S} : d(P_\theta, P_0) > T\varepsilon_N} \prod_{i=1}^N \frac{p_\theta(Y_i|X_i)}{p_0(Y_i|X_i)} Pr(\mathrm{d}\theta)(1 - \phi_N) \\
\leq \; & \Pi(\mathcal{S} \setminus \mathcal{S}_N) + \int_{\theta \in \mathcal{S}_N : d(P_\theta, P_0) > T\varepsilon_N} P_\theta^N (1 - \phi_N) Pr(\mathrm{d}\theta) \\
\leq \; & \Pi(\mathcal{S} \setminus \mathcal{S}_N) + e^{-LNT^2\varepsilon_N^2} \quad \text{by (12)} \\
\leq \; & 2e^{-N\varepsilon_N^2(C+4)}
\end{aligned} \tag{14}$$

if $KM^2 \geq C + 4$, by Assumption 2.

By Lemma 1 (stated below) and Assumption 3, with probability at least $1-1/(C^2N\varepsilon_N^2)$, we have

$$\int \prod_{i=1}^{N} \frac{p_\theta(Y_i|X_i)}{p_0(Y_i|X_i)} Pr(\mathrm{d}\theta) \geq e^{-2N\varepsilon_N^2} Pr(B_n) \geq e^{-N\varepsilon_N^2(2+C)} \tag{15}$$

where

$$B_n = \left\{ \theta : -P_0 \log \frac{p_\theta(Y_1|X_1)}{p_0(Y_1|X_1)} \leq \varepsilon_N^2,\ P_0 \left(\frac{p_\theta(Y_1|X_1)}{p_0(Y_1|X_1)} \right)^2 \leq \varepsilon_N^2 \right\}$$

Let A_N be the event that (15) holds. We have

$$\begin{aligned}
& P_0^N Pr(\theta : d(P_\theta, P_0) > T\varepsilon_N | X, Y)(1-\phi_N)\mathbf{1}_{A_N} \\
= \ & P_0^N \frac{\int_{\theta: d(P_\theta,P_0)>T\varepsilon_N} \prod_{i=1}^N \frac{p_\theta(Y_i|X_i)}{p_0(Y_i|X_i)} Pr(\mathrm{d}\theta)}{\int \prod_{i=1}^N \frac{p_\theta(Y_i|X_i)}{p_0(Y_i|X_i)} Pr(\mathrm{d}\theta)} (1-\phi_N)\mathbf{1}_{A_N} \\
\leq \ & e^{N\varepsilon_N^2(2+C)} 2e^{-N\varepsilon_N^2(C+4)} \ \text{ by (14) and (15)} \\
= \ & 2e^{-2N\varepsilon_N^2}
\end{aligned}$$

Therefore,

$$\begin{aligned}
& P_0^N Pr(\theta : d(P_\theta, P_0) > T\varepsilon_N | X, Y) \\
= \ & P_0^N Pr(\theta : d(P_\theta, P_0) > T\varepsilon_N | X, Y)\phi_N + P_0^N Pr(\theta : d(P_\theta, P_0) > T\varepsilon_N | X, Y)(1-\phi_N)\mathbf{1}_{A_N} \\
& + P_0^N Pr(\theta : d(P_\theta, P_0) > T\varepsilon_N | X, Y)(1-\phi_N)(1-\mathbf{1}_{A_N}) \\
\leq \ & P_0^N Pr(\theta : d(P_\theta, P_0) > T\varepsilon_N | X, Y)\phi_N + P_0^N Pr(\theta : d(P_\theta, P_0) > T\varepsilon_N | X, Y)(1-\phi_N)\mathbf{1}_{A_N} \\
& + P_0^N(A_N^c) \ \text{ for sufficiently large } T \\
\leq \ & 2e^{-LN\varepsilon_N^2} + 2e^{-2N\varepsilon_N^2} + \frac{1}{C^2 N\varepsilon_N^2}
\end{aligned}$$

by (13), (16) and the property of A_N. By Chebyshev's inequality, we have

$$P_0^N \left(Pr(\theta : d(P_\theta, P_0) > T\varepsilon_N^2 | X, Y) > \delta \right) \leq \frac{1}{C^2 N\varepsilon_N^2 \delta} + \frac{2e^{-LN\varepsilon_N^2}}{\delta} + \frac{2e^{-2N\varepsilon_N^2}}{\delta}$$

which concludes the theorem.

Lemma 1 *For any $\epsilon > 0$ and probability distribution Π defined on the set*

$$\left\{ \theta : -P_0 \log \frac{p_\theta(Y|X)}{p_0(Y_1|X_1)} \leq \epsilon^2,\ P_0 \left(\frac{p_\theta(Y|X)}{p_0(Y_1|X_1)} \right)^2 \leq \epsilon^2 \right\} \tag{16}$$

we have, for every $C > 0$,

$$P_0^N \left(\int \prod_{i=1}^{N} \frac{p_\theta(Y_i|X_i)}{p_0(Y_i|X_i)} \Pi(\mathrm{d}\theta) \leq e^{-(1+C)N\epsilon^2} \right) \leq \frac{1}{C^2 N\epsilon^2} \tag{17}$$

The proof follows closely with Lemma 8.1 in [Ghosal et al., 2000].

Proof. By Jensen's inequality, we have

$$\log \int \prod_{i=1}^{N} \frac{p_\theta(Y_i|X_i)}{p_0(Y_i|X_i)} \Pi(\mathrm{d}\theta) \geq \int \sum_{i=1}^{N} \log \frac{p_\theta(Y_i|X_i)}{p_0(Y_i|X_i)} \Pi(\mathrm{d}\theta)$$

Hence the left hand side of (17) is bounded from above by

$$P_0^N \left(\mathcal{G}_N \int \log \frac{p_\theta(Y|X)}{p_0(Y_1|X_1)} \Pi(\mathrm{d}\theta) \leq -(1+C)\sqrt{N}\epsilon^2 - \sqrt{N} P_0 \int \log \frac{p_\theta(Y_1|X_1)}{p_0(Y_1|X_1)} \Pi(\mathrm{d}\theta) \right) \quad (18)$$

where $\mathcal{G}_N = \sqrt{n}(P_N - P_0)$ denotes the empirical process generated by (X, Y), and P_N here is the empirical distribution. Since Π is defined on (16), by Fubini's theorem, (18) is bounded further by

$$P_0^N \left(\mathcal{G}_N \int \log \frac{p_\theta(Y|X)}{p_0(Y|X)} \Pi(\mathrm{d}\theta) \leq -C\sqrt{N}\epsilon^2 \right) \quad (19)$$

By Chebyshev's inequality, (19) is bounded by

$$\frac{Var_0 \left(\int \log \frac{p_\theta(Y_1|X_1)}{p_0(Y_1|X_1)} \Pi(\mathrm{d}\theta) \right)}{C^2 N \epsilon^4}$$

where Var_0 denotes the variance under P_0, which is in turn bounded by

$$\frac{P_0 \left(\int \log \frac{p_\theta(Y_1|X_1)}{p_0(Y_1|X_1)} \Pi(\mathrm{d}\theta) \right)^2}{C^2 N \epsilon^4} \leq \frac{\int P_0 \left(\log \frac{p_\theta(Y_1|X_1)}{p_0(Y_1|X_1)} \right)^2 \Pi(\mathrm{d}\theta)}{C^2 N \epsilon^4}$$

by Jensen's inequality and Fubini's theorem again. The last expression is bounded by $1/(C^2 N \epsilon^2)$ since Π is defined on (16).

Proof of Theorem 2

Proof of 1. Consider $N \to \infty$ and fix a $T > 0$. We have

$$\begin{aligned}
& Pr(\theta : d(P_\theta, P_0) \leq T\varepsilon_N^2 | X, Y) \\
= \; & E_{Pr} \left[Pr(\theta : d(P_\theta, P_0) \leq T\varepsilon_N^2 | M_k, X, Y) | X, Y \right] \quad (20) \\
& \quad \text{where } E_{Pr}[\cdot|X, Y] \text{ denotes the posterior expectation} \\
& \quad \text{and } Pr(\cdot|M_k, X, Y) \text{ denotes the posterior distribution given model } M_k \\
= \; & Pr(M_1|X, Y) Pr(\theta : d(P_\theta, P_0) \leq T\varepsilon_N^2 | M_1, X, Y) \quad (21)
\end{aligned}$$

by the condition that $d(P_\theta, P_0) > M\varepsilon_N^2$ for any $\theta \in \mathcal{S}_{-1}$ eventually. Hence

$$Pr(\theta : d(P_\theta, P_0) \leq M\varepsilon_N^2 | X, Y) \geq 1 - \delta \quad (22)$$

implies

$$Pr(M_1|X, Y) \geq 1 - \delta \quad (23)$$

The result then follows from Theorem 1, which implies that (22) occurs with probability at least

$$1 - \left(\frac{1}{C^2 N \varepsilon_N^2 \delta} + \frac{2e^{-LN\varepsilon_N^2}}{\delta} + \frac{2e^{-2N\varepsilon_N^2}}{\delta} \right)$$

$$P_0^N \left(Pr(M_1|X,Y) < 1 - \delta \right) \leq \frac{1}{C^2 N \varepsilon_N^2 \delta} + \frac{2e^{-LN\varepsilon_N^2}}{\delta} + \frac{2e^{-2N\varepsilon_N^2}}{\delta} \tag{24}$$

Proof of 2. Note that (23) implies

$$d_E(Pr(M_k|X,Y), \mathbf{e}_1) = \sqrt{(1 - Pr(M_1|X,Y))^2 + \sum_{k \neq 1} Pr(M_k|X,Y)^2} \leq \sqrt{2}\delta \tag{25}$$

since $(1 - Pr(M_1|X,Y))^2 \leq \delta^2$ and $(\delta, 0, \ldots, 0)$ is an optimizer of the optimization

$$\max \sum_{i=2}^{K} x_i^2 \;\text{ subject to }\; \sum_{i=2}^{K} x_i \leq \delta$$

Hence (24) and (25) together imply

$$P_0^N \left(d_E(Pr(M_k|X,Y), \mathbf{e}_1) \geq \sqrt{2}\delta \right) \leq \frac{1}{C^2 N \varepsilon_N^2 \delta} + \frac{2e^{-LN\varepsilon_N^2}}{\delta} + \frac{2e^{-2N\varepsilon_N^2}}{\delta}$$

By redefining $\tilde{\delta} = \sqrt{2}\delta$, we get (6).

Proof of 3. Note that (23) implies

$$\begin{aligned} d_H(Pr(M_k|X,Y), \mathbf{e}_1) &= \sqrt{\frac{1}{2}\left((\sqrt{1 - Pr(M_1|X,Y)})^2 + \sum_{k \neq 1} Pr(M_k|X,Y) \right)} \\ &\leq \sqrt{\frac{1}{2}\left((1 - \sqrt{1-\delta})^2 + \delta \right)} \end{aligned} \tag{26}$$

since $x_i = \delta/(k-1)$ for all $i \neq 0$ gives the optimizer of the optimization

$$\max \sum_{i \neq 0} \sqrt{x_i} \;\text{ subject to }\; \sum_{i \neq 0} x_i \leq \delta$$

Hence (24) and (26) together imply

$$P_0^N \left(d_H(Pr(M_k|X,Y), \mathbf{e}_1) > \sqrt{\frac{1}{2}\left((1 - \sqrt{1-\delta})^2 + \delta \right)} \right) \leq \frac{1}{C^2 N \varepsilon_N^2 \delta} + \frac{2e^{-LN\varepsilon_N^2}}{\delta} + \frac{2e^{-2N\varepsilon_N^2}}{\delta} \tag{27}$$

Note that $(1 - \sqrt{1-\delta})^2$ is a convex function in δ for $0 < \delta < 1$ and is equal to 0 at $\delta = 0$. Thus $(1 - \sqrt{1-\delta})^2 \leq (\sqrt{2}-1)^2\delta$ for $0 < \delta < 1/2$, where $(\sqrt{2}-1)^2$ is the slope of the line between $(0,0)$ and $(1/2, (1 - \sqrt{1-1/2})^2$. Hence, for $0 < \delta < 1/2$, we have

$$\sqrt{\frac{1}{2}\left((1 - \sqrt{1-\delta})^2 + \delta \right)} \leq \sqrt{((\sqrt{2}-1)^2 + 1)\frac{\delta}{2}}$$

Combining with (27), we have

$$P_0^N\left(d_H(Pr(M_k|X,Y),\mathbf{e}_1) > \sqrt{((\sqrt{2}-1)^2+1)\frac{\delta}{2}}\right) \le \frac{1}{C^2 N\varepsilon_N^2\delta} + \frac{2e^{-LN\varepsilon_N^2}}{\delta} + \frac{2e^{-2N\varepsilon_N^2}}{\delta}$$

By redefining $\tilde{\delta} = \sqrt{((\sqrt{2}-1)^2+1)\delta/2}$, we get (7).

Remark 1 *Note that the concentration bound for Hellinger distance* (7) *is inferior to that for Euclidean distance* (6) *for small* δ *since* δ^2 *instead of* δ *shows up in the RHS of* (7)*. This is because in our proof, the function* $\sqrt{(1-\sqrt{1-\delta})^2+\delta}$ *that appears in* (27) *has derivative* $1/(2\sqrt{(1-\delta)((1-\sqrt{1-\delta})^2+\delta)})$ *which is* ∞ *at* $\delta = 0$*, and thus no linearization is available when* δ *is close to 0. In this sense the bound in* (7) *cannot be improved.*

Proof of Theorem 3 and 4

The proofs of both theorems rely on a key theorem on geometric median in [Minsker, 2015], restated below:

Theorem 5 *Consider a Hilbert space* $(\mathbb{H}, \langle\cdot,\cdot\rangle)$ *and* $\xi_0 \in \mathbb{H}$*. Let* $\hat{\xi}_1,\dots,\hat{\xi}_R \in \mathbb{H}$ *be a collection of independent random* $\mathbb{H}$*-valued elements. Let* α, q, ν *be constants such that* $0 < q < \alpha < 1/2$ *and* $0 \le \nu < (\alpha - q)/(1-q)$*. Suppose that there exists* $\epsilon > 0$ *such that for all* j*, where* $1 \le j \le \lfloor(1-\nu)R\rfloor + 1$*,*

$$P(\|\hat{\xi}_j - \xi_0\| > \epsilon) \le q$$

Let $\hat{\xi}_* = med_g(\hat{\xi}_1,\dots,\hat{\xi}_R)$ *be the geometric median of* $\{\hat{\xi}_1,\dots,\hat{\xi}_R\}$*. Then*

$$P(\|\hat{\xi}_* - \xi_0\| > C_\alpha\epsilon) \le \left(e^{(1-\nu)\psi(\frac{\alpha-\nu}{1-\nu},q)}\right)^{-R}$$

where $C_\alpha = (1-\alpha)\sqrt{1/(1-2\alpha)}$*, and*

$$\psi(\alpha,q) = (1-\alpha)\log\frac{1-\alpha}{1-q} + \alpha\log\frac{\alpha}{q}$$

We focus on Theorem 3. The proof for Theorem 4 is a straightforward modification in light of Theorem 5.

Proof of 1. Immediate by noting that

$$P_0^s\left(d(Pr(M_k|X_{(j)},Y_{(j)}),\mathbf{e}_1) > \delta\right) \le q$$

for all $j = 1,\dots,R$, and applying Theorem 5.

Proof of 2. Note that

$$d_E(Pr_*(M_k|X,Y),\mathbf{e}_1) \ge (1 - Pr_*(M_1|X,Y))\sqrt{\frac{K}{K-1}} \tag{28}$$

To see this, let $a = Pr_*(M_1|X,Y)$. We have $d_E(Pr_*(M_k|X,Y), \mathbf{e}_1) = \sqrt{(1-a)^2 + \sum_{i=2}^K x_i^2}$ where x_i's satisfy $\sum_{i=2}^K x_i = 1-a$. Since $(1-a)/(K-1)$ is the optimizer of the optimization

$$\min \sum_{i=2}^K x_i^2 \ \text{ subject to } \ \sum_{i=2}^K x_i = 1-a$$

we get $\sqrt{(1-a)^2 + \sum_{i=2}^K x_i^2} \geq (1-a)\sqrt{K/(K-1)}$.

Hence (1) and (28) together give

$$P_0^N \left(Pr_*(M_1|X,Y) < 1 - C_\alpha \delta \sqrt{\frac{K-1}{K}} \right) \leq \left(e^{(1-\nu)\psi(\frac{\alpha-\nu}{1-\nu}, q)} \right)^{-R}$$

Proof of 3. By using the same thread of argument as in (21), we have

$$Pr(\theta : d(P_\theta, P_0) \leq T\varepsilon_N^2 | X, Y) \geq 1 - \delta$$

implies

$$Pr(\theta : d(P_\theta, P_0) \leq T\varepsilon_N^2 | M_1, X, Y) \geq 1 - \delta$$

for any $0 < \delta < 1$. Hence (4) implies that

$$P_0^N \left(Pr(\theta : d(P_\theta, P_0) > T\varepsilon_N^2 | M_1, X, Y) > \delta \right) \leq \frac{1}{C^2 N \varepsilon_N^2 \delta} + \frac{2e^{-LN\varepsilon_N^2}}{\delta} + \frac{2e^{-2N\varepsilon_N^2}}{\delta}$$

and by the additional assumption we have

$$P_0^N \left(Pr(\theta : \tilde{C} d_{\mathcal{F}}(P_\theta, P_0)^\gamma > T\varepsilon_N^2 | M_1, X, Y) > \delta \right) \leq \frac{1}{C^2 N \varepsilon_N^2 \delta} + \frac{2e^{-LN\varepsilon_N^2}}{\delta} + \frac{2e^{-2N\varepsilon_N^2}}{\delta}$$

Applying a union bound, we get

$$\begin{aligned} & P_0^N \Bigg(Pr_*(M_1|X,Y) > 1 - C_\alpha \delta \sqrt{\frac{K-1}{K}}, \\ & \quad Pr_* \left(\theta : d_{\mathcal{F}}(P_\theta, P_0) \leq \left(\frac{T}{\tilde{C}} \varepsilon_N^2 \right)^{1/\gamma} \middle| M_1, X, Y \right) \geq 1 - \delta \Bigg) \\ \geq \ & 1 - 2 \left(e^{(1-\nu)\psi(\frac{\alpha-\nu}{1-\nu}, q)} \right)^{-R} \end{aligned}$$

Gibbs Sampler for Parallel MCMC for Spike and Slab Prior

For posterior inference in the spike and slab model, let $\Delta = \text{diag}\left\{J_1\tau_1^2, \ldots, J_D\tau_D^2\right\}$, we can perform Gibbs sampling by drawing from the following posteriors:

$$Pr(\beta_{SS}|-) \sim N(\mu_{\beta_{SS}}, \Sigma_{\beta_{SS}}), \tag{29}$$

$$\Sigma_{\beta_{SS}} = \left(\Delta^{-1} + \frac{P}{N\sigma^2}X^TX\right)^{-1},\ \mu_{\beta_{SS}} = \Sigma_{\beta_{SS}}\left(\frac{P}{N\sigma^2}X^TY\right), \tag{30}$$

$$Pr\left(\sigma_{SS}^{-2}|-\right) \sim \text{Gamma}\left(a + \frac{N}{2}, b + \frac{P}{2N}\left(Y - X\beta\right)^T\left(Y - X\beta\right)\right) \tag{31}$$

$$\begin{aligned} Pr\left(J_d|-\right) &\propto w_{d1}\delta_{J_d}(\nu_0) + w_{d2}\delta_{J_d}(1) \\ w_{d1} &= (1-w)\nu_0^{-1/2}\exp\left\{-\frac{\beta_d^2}{2\nu_0\tau_d^2}\right\} \\ w_{d2} &= w\exp\left\{-\frac{\beta_d^2}{2\tau_d^2}\right\} \end{aligned} \tag{32}$$

$$Pr\left(\tau_d^{-2}|-\right) \sim \text{Gamma}\left(a_\tau + \frac{1}{2}, b_\tau + \frac{\beta_d^2}{2J_d}\right) \tag{33}$$

$$Pr(w|-) \sim \text{Beta}\left(1 + |\{d : J_d = 1\}|, 1 + |\{d : J_d = \nu_0\}|\right) \tag{34}$$

Additional Results and Plots From Simulation Studies

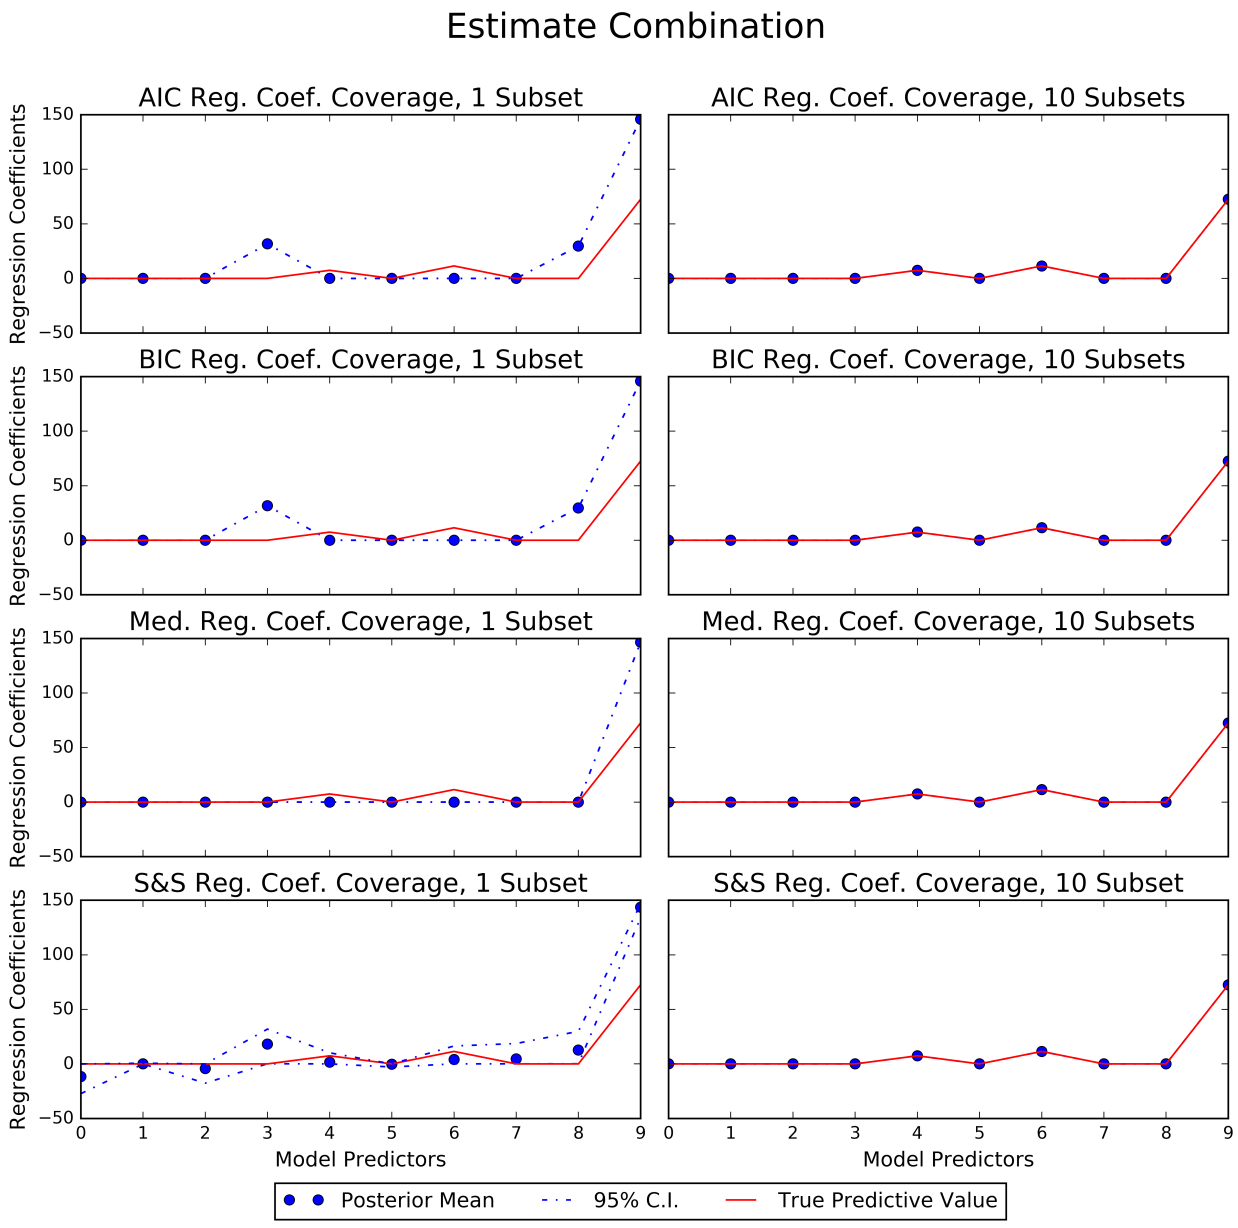
Estimate Combination
AIC Reg. Coef. Coverage, 1 Subset
AIC Reg. Coef. Coverage, 10 Subsets
BIC Reg. Coef. Coverage, 1 Subset
BIC Reg. Coef. Coverage, 10 Subsets
Med. Reg. Coef. Coverage, 1 Subset
Med. Reg. Coef. Coverage, 10 Subsets
S&S Reg. Coef. Coverage, 1 Subset
S&S Reg. Coef. Coverage, 10 Subset
Regression Coefficients
Model Predictors
Posterior Mean
95% C.I.
True Predictive Value

Figure 4: Posterior regression parameter coverage test results, estimate combination.

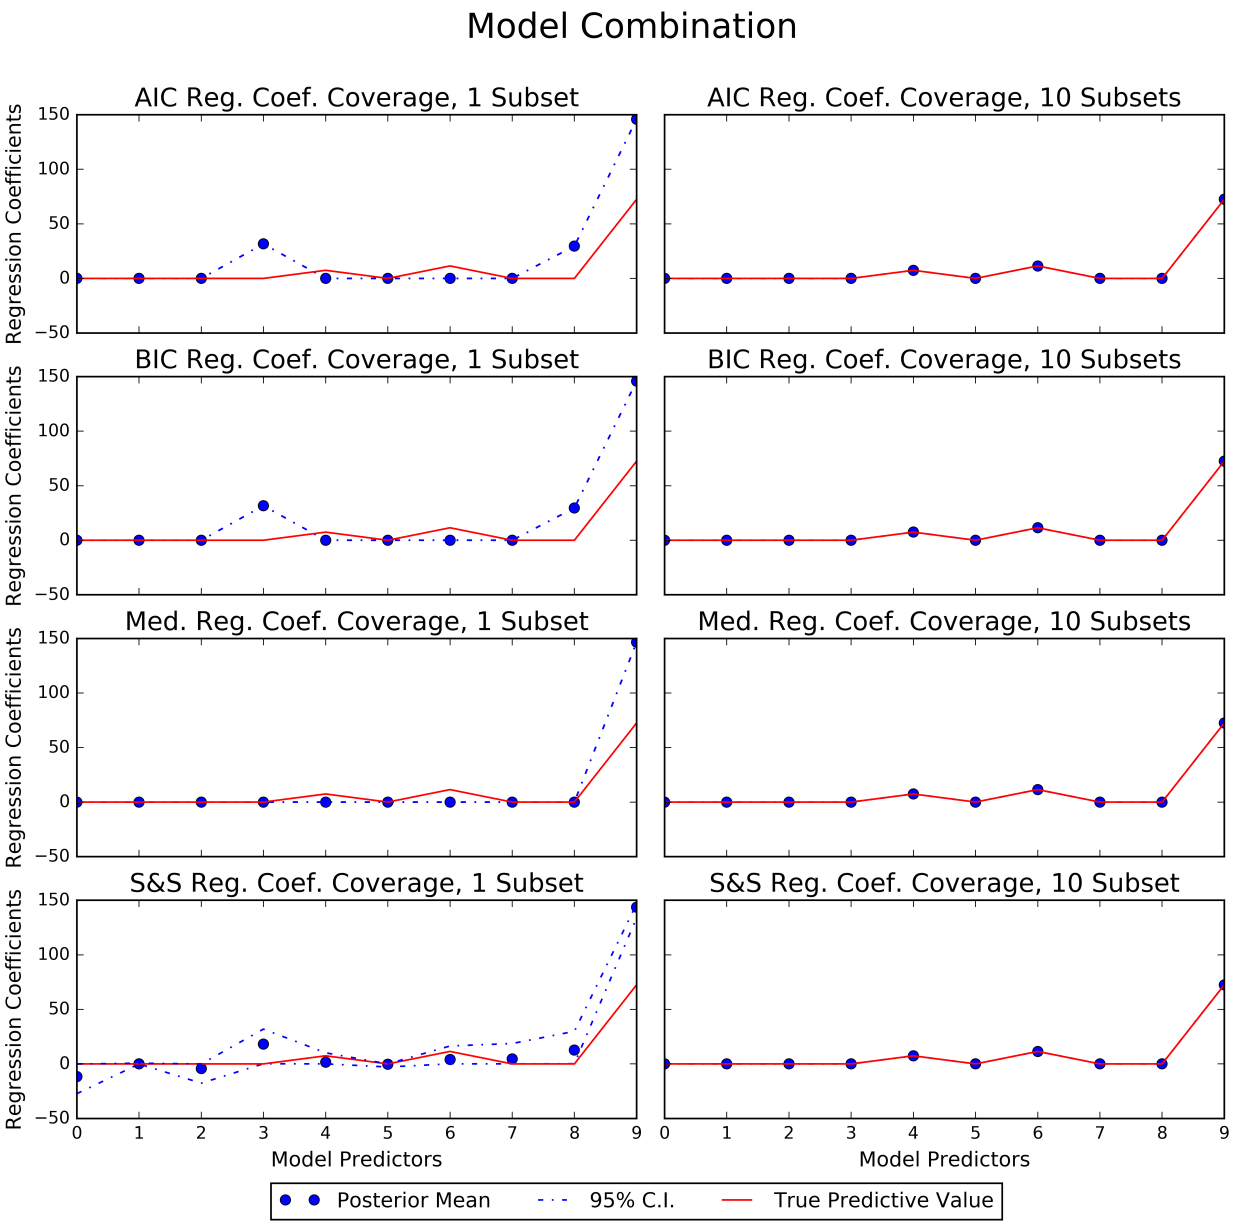
Model Combination
AIC Reg. Coef. Coverage, 1 Subset
AIC Reg. Coef. Coverage, 10 Subsets
BIC Reg. Coef. Coverage, 1 Subset
BIC Reg. Coef. Coverage, 10 Subsets
Med. Reg. Coef. Coverage, 1 Subset
Med. Reg. Coef. Coverage, 10 Subsets
S&S Reg. Coef. Coverage, 1 Subset
S&S Reg. Coef. Coverage, 10 Subset
Regression Coefficients
Model Predictors
Posterior Mean
95% C.I.
True Predictive Value

Figure 5: Posterior regression parameter coverage test results, model combination.

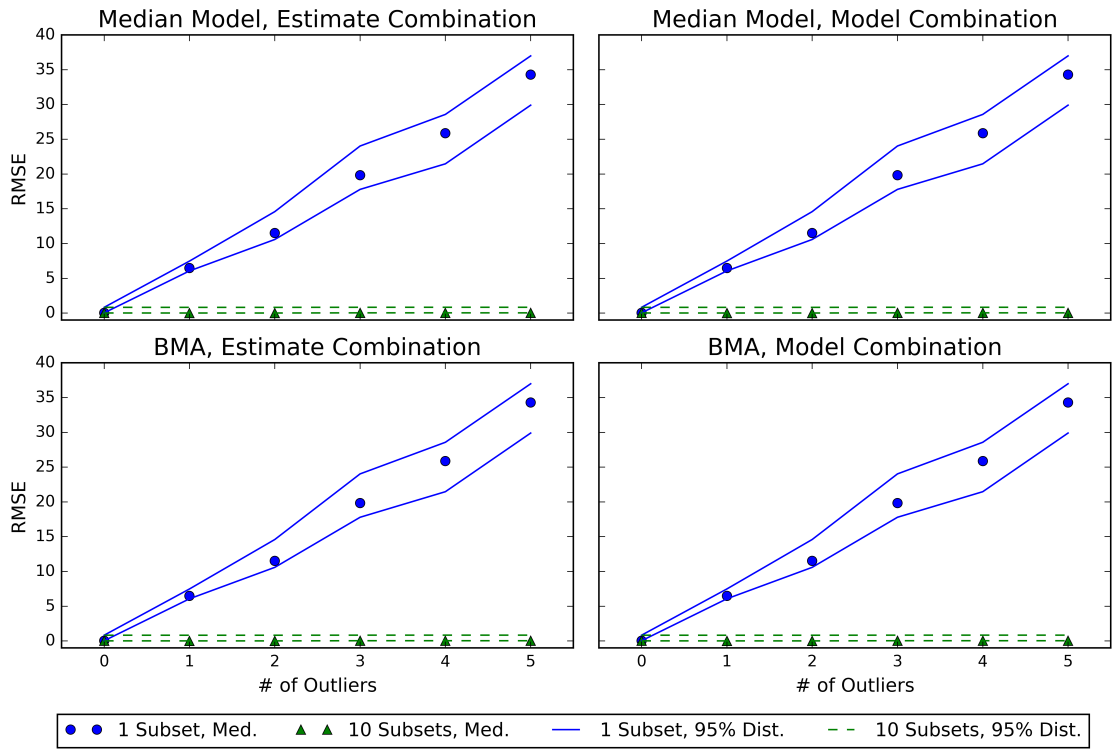

Figure 6: Contamination test

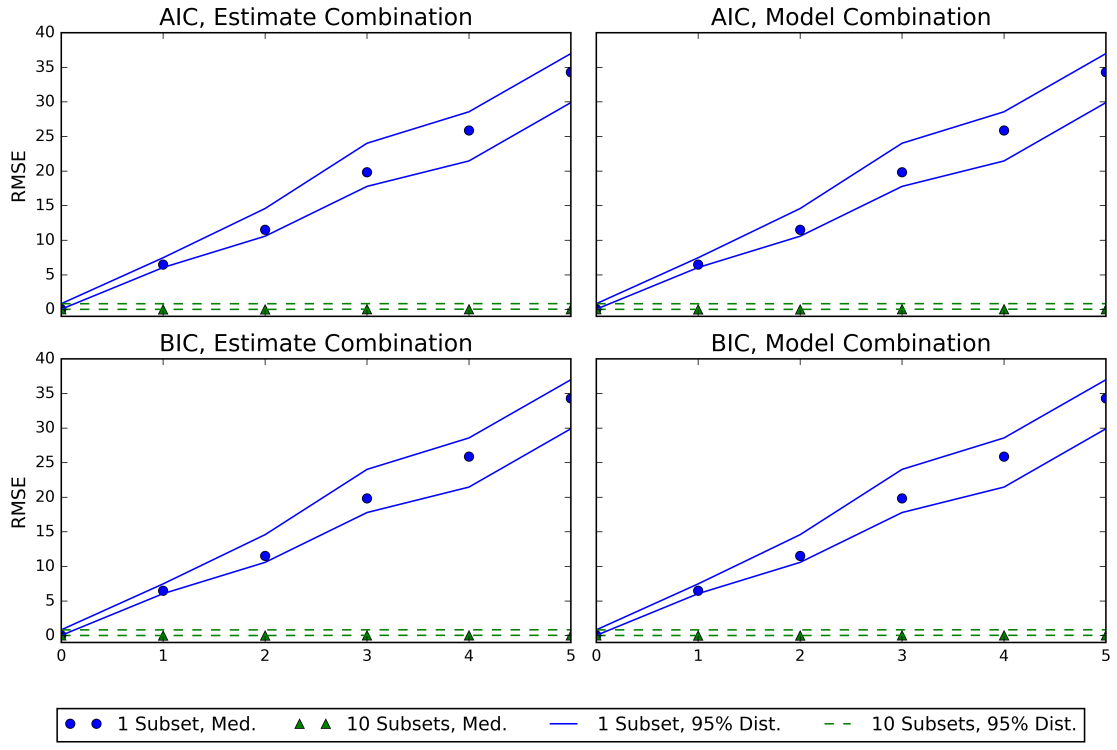

Figure 7: Contamination test

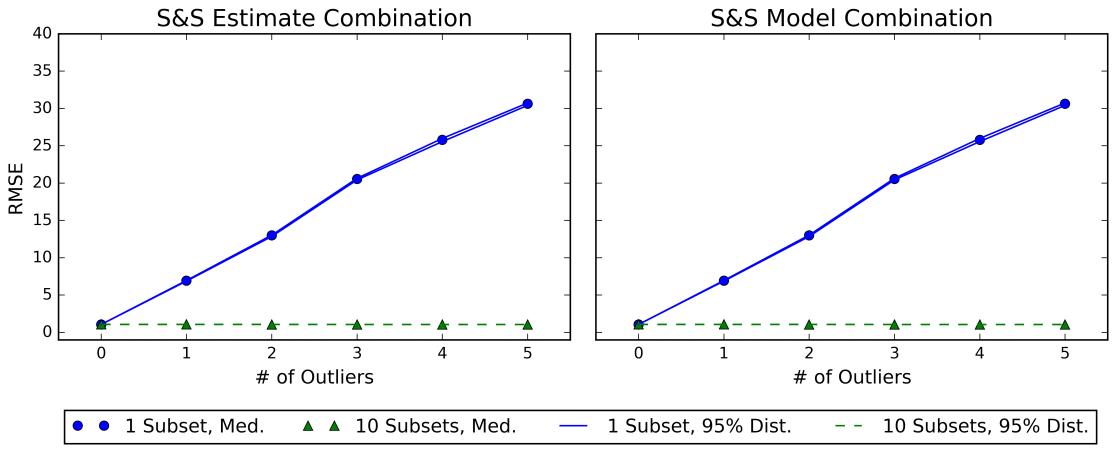

Figure 8: Contamination test.

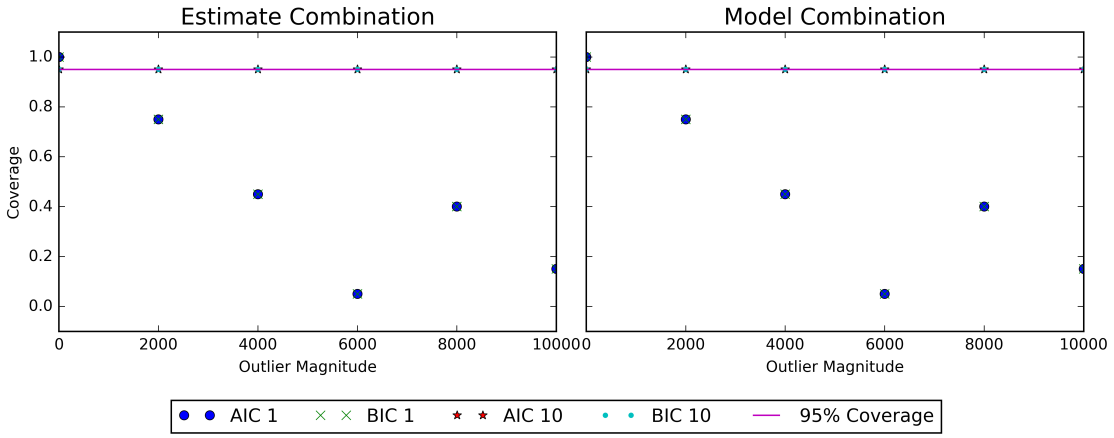

Figure 9: Coverage test.

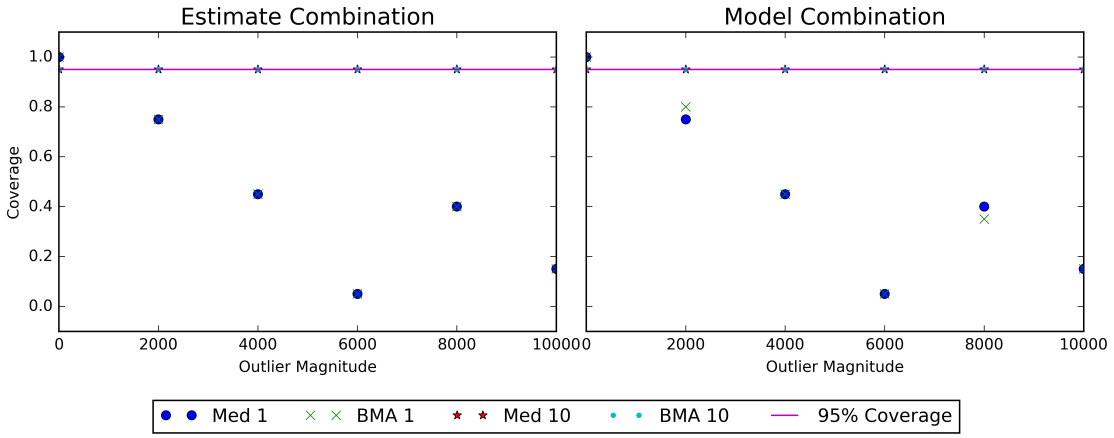

Figure 10: Coverage test.

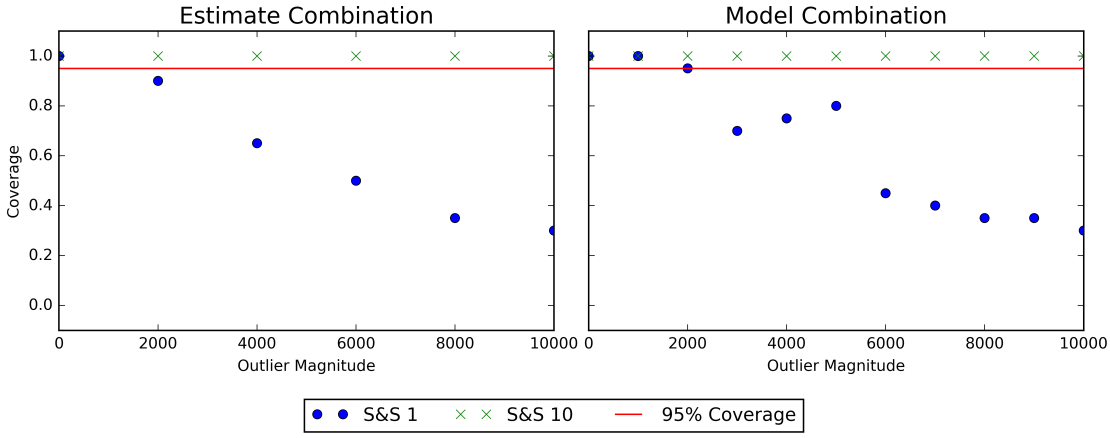

Figure 11: Coverage test.

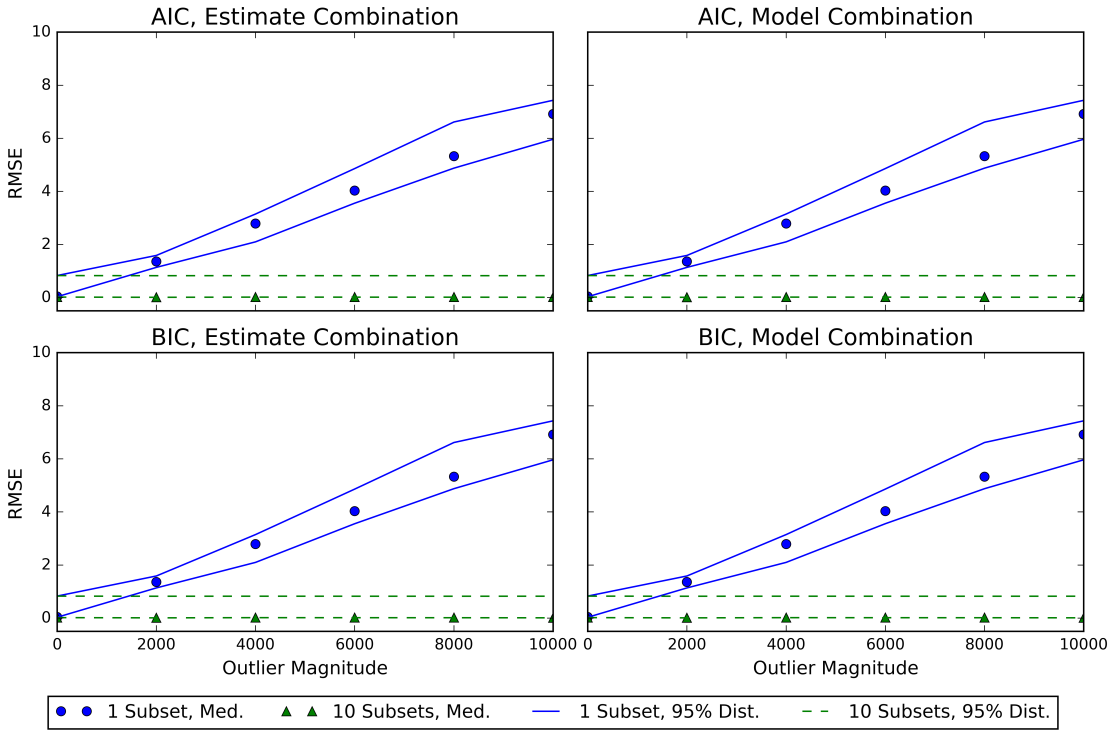

Figure 12: Magnitude test.

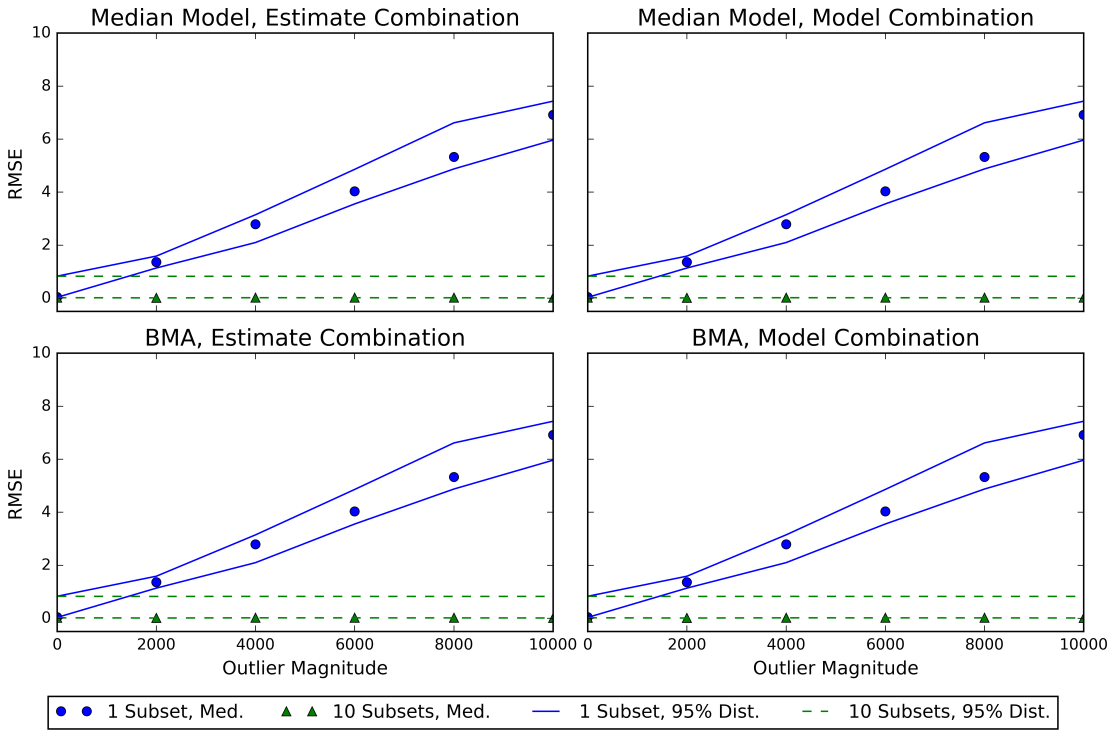

Figure 13: Magnitude test.

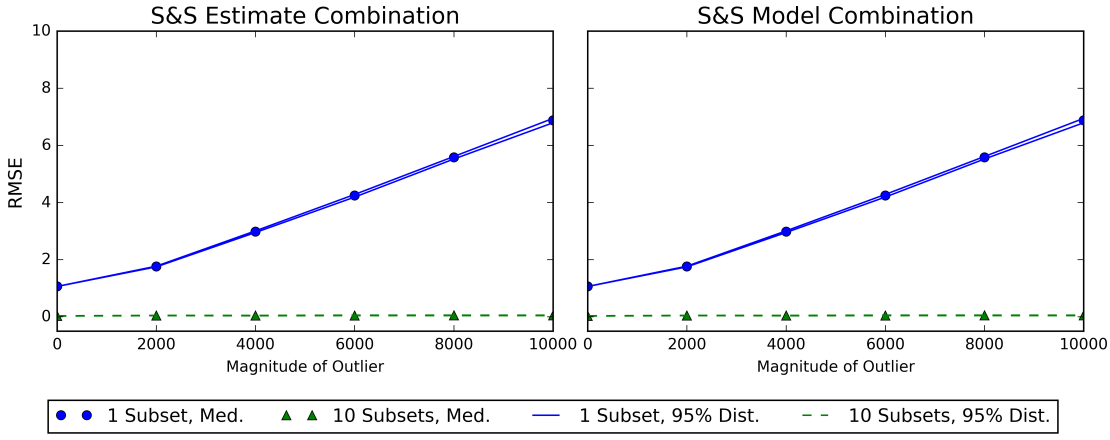

Figure 14: Magnitude test.